



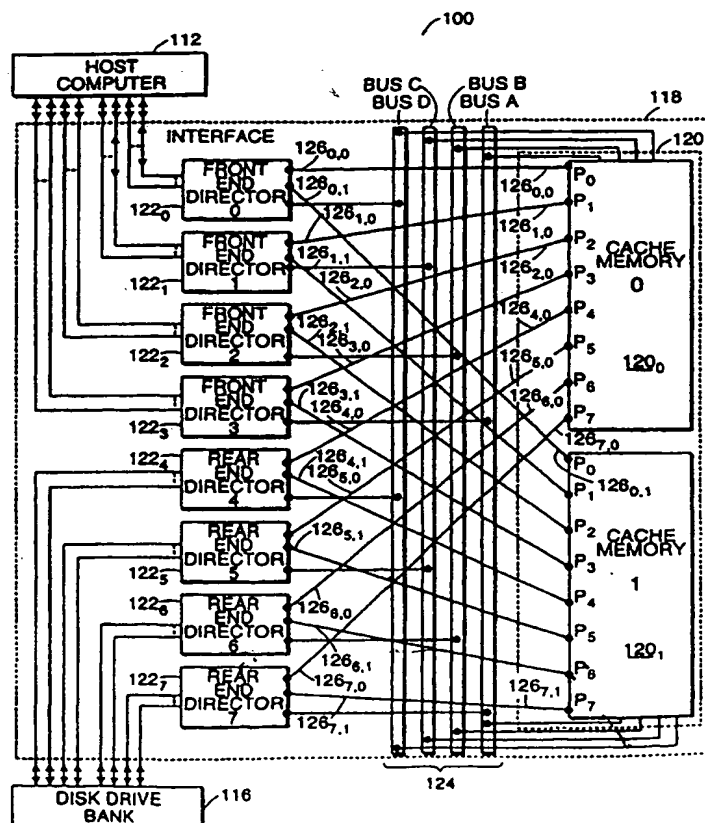
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : G06F 13/12, 12/08, 3/06, 13/40	A1	(11) International Publication Number: WO 00/39691 (43) International Publication Date: 6 July 2000 (06.07.00)
(21) International Application Number: PCT/US99/31213 (22) International Filing Date: 29 December 1999 (29.12.99) (30) Priority Data: 09/223,519 30 December 1998 (30.12.98) US (71) Applicant: EMC CORPORATION [US/US]; 171 South Street, Hopkinton, MA 01748-9103 (US). (72) Inventors: WALTON, John, K.; 22 King Philip Path, Mendon, MA 01756 (US). CASTEL, Daniel; 6 McCarthy Circle, Framingham, MA 01702 (US). CHILTON, Kendell, Alan; 53 Alan Road, Marlborough, MA 01752 (US). (74) Agent: SHARKANSKY, Richard, M.; Fish & Richardson P.C., 225 Franklin Street, Boston, MA 02110-2804 (US).		(81) Designated States: JP, KR, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>

(54) Title: DATA STORAGE SYSTEM

(57) Abstract

A data storage system wherein end-user data is transferred between a host computer and a bank of disk drives through an interface. The interface includes a memory and a plurality of directors interconnected through an interface state data bus and a plurality of end-user data busses. At least one front-end one of the directors is in communication with the host computer and at least one rear-end one of the directors is in communication with the bank of disk drives. The interface state data bus section is in communication with: both the at least one front-end one and the at least one rear-end one of the directors; and to the memory. Each one of the plurality of end-user data busses has a first end coupled to a corresponding one of the plurality of directors and a second end coupled to the memory. The plurality of directors control the end-user data transfer between the host computer and the bank of disk drives through the memory in response to interface state data generated by the directors as such end-user data passes through the end-user data busses. The generated interface state data is transferred among the directors through the memory as such end-user data passes through the end-user bus.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

- 1 -

DATA STORAGE SYSTEMBackground of the Invention

5 This invention relates generally to data storage systems, and more particularly to data storage systems having redundancy arrangements to protect against total system failure in the event of a failure in a component or subassembly of the storage system.

10 As is known in the art, large host computer systems require large capacity data storage systems. These large computer systems generally includes data processors which perform many operations on data introduced to the computer system through peripherals
15 including the data storage system. The results of these operations are output to peripherals, including the storage system.

One type of data storage system is a magnetic disk storage system. Here a bank of disk drives and the
20 computer system are coupled together through an interface. The interface includes CPU, or "front end", directors (or controllers) and "back end" disk directors (or controllers). The interface operates the directors in such a way that they are transparent to the computer.
25 That is, data is stored in, and retrieved from, the bank of disk drives in such a way that the computer system merely thinks it is operating with one large memory. One such system is described in U.S. Patent 5,206,939, entitled "System and Method for Disk Mapping and Data
30 Retrieval", inventors Moshe Yanai, Natan Vishlitzky, Bruno Alterescu and Daniel Castel, issued April 27, 1993, and assigned to the same assignee as the present invention.

As described in such U.S. Patent, the interface
35 may also include, in addition to the CPU directors and disk directors, addressable cache memories. The cache

- 2 -

memory is a semiconductor memory and is provided to rapidly store data from the computer system before storage in the disk drives, and, on the other hand, store data from the disk drives prior to being sent to the
5 computer. The cache memory being a semiconductor memory, as distinguished from a magnetic memory as in the case of the disk drives, is much faster than the disk drives in reading and writing data.

The CPU directors, disk directors and cache memory
10 are interconnected through a backplane printed circuit board. More particularly, disk directors are mounted on disk director printed circuit boards. CPU directors are mounted on CPU director printed circuit boards. And, cache memories are mounted on cache memory printed
15 circuit boards. The disk director, CPU director and cache memory printed circuit boards plug into the backplane printed circuit board. In order to provide data integrity in case of a failure in a director, the backplane printed circuit board has a pair of buses. One
20 set of the disk directors is connected to one bus and another set of the disk directors is connected to the other bus. Likewise, one set the CPU directors is connected to one bus and another set of the CPU directors is connected to the other bus. The cache memories are
25 connected to both buses. Each one of the buses provides data, address and control information. Thus, the use of two buses provides a degree of redundancy to protect against a total system failure in the event that the directors, or disk drives connected to one bus fail and
30 also increases the bandwidth of the system compared with a system which uses a single bus. One such dual bus system is shown in FIG. 1.

In operation, when the host computer wishes to store end-user (i.e., host computer) data at an address,
35 the host computer issues a write request to one of the

- 3 -

front-end directors to perform a write command. One of the front-end directors replies to the request and asks the host computer for the data. After the request has passed to the requesting one of the front-end directors, 5 the director determines the size of the end-user data and reserves space in the cache memory to store the request. The front-end director then produces control signals on either one of the busses connected to such front-end director. The host computer then transfers the data to 10 the front-end director. The front-end director then advises the host computer that the transfer is complete. The front-end director looks up in a Table, not shown, stored in the cache memory to determine which one of the rear-end directors is to handle this request. The Table 15 maps the host computer address into an address in the bank of disk drives. The front-end director then puts a notification in a "mail box" (not shown and stored in the cache memory) for the rear-end director which is to handle the request, the amount of the data and the disk 20 address for the data. Other rear-end directors poll the cache memory when they are idle to check their "mail boxes". If the polled "mail box" indicates a transfer is to be made, the rear-end director processes the request, addresses the disk drive in the bank, reads the data from 25 the cache memory and writes it into the addresses of a disk drive in the bank. When end-user data previously stored in the bank of disk drives is to be read from the disk drive and returned to the host computer, the interface system operates in a reciprocal manner. The 30 internal operation of the interface, (e.g. "mail-box polling", event flags, data structures, device tables, queues, etc.) is controlled by interface state data which passes between the directors through the cache memory. Further, end-user data is transferred through the 35 interface as a series of multi-word transfers, or bursts.

- 4 -

Each word transfer in a multi-word transfer is here, for example, 64 bits. Here, an end-user data transfer is made up of, for example, 32 bursts. Each interface state word is a single word having, for example, 64 bits.

5 It is first noted that the end-user data and interface state data are transferred among the directors and the cache memory on the busses. The transfer of each word, whether a burst of end-user data or an interface state data passes through the interface in the same
10 manner; i.e., requiring a fixed amount of overhead, i.e., bus arbitration, etc. Each one of the two busses must share its bandwidth with both end-user data and the interface state data. Therefore, the bandwidth of the system may not be totally allocated to end-user data
15 transfer between the host computer and the bank of disk drives.

Summary of the Invention

In accordance with the present invention, a data storage system is provided wherein end-user data is
20 transferred between a host computer and a bank of disk drives through an interface. The interface includes a memory and a plurality of directors interconnected through an interface state data bus and a plurality of end-user data busses. At least one front-end one of the
25 directors is in communication with the host computer and at least one rear-end one of the directors is in communication with the bank of disk drives. The interface state data bus section is in communication with: both the at least one front-end one and the at
30 least one rear-end one of the directors; and to the memory. Each one of the plurality of end-user data buses has a first end coupled to a corresponding one of the plurality of directors and a second end coupled to the memory. The plurality of directors control the end-user
35 data transfer between the host computer and the bank of

- 5 -

disk drives through the memory in response to interface state data generated by the directors as such end-user data passes through the end-user data busses. The generated interface state data is transferred among the directors through the memory as such end-user data passes through the end-user bus.

With such an arrangement, the system bandwidth is increased because end-user data and interface state data are carried on separate bus systems within the interface.

10 Brief Description of the Drawing

For a more complete understanding of the invention, reference is now made to the following description taken together in conjunction with the accompanying drawing, in which:

15 FIG. 1 is a block diagram of a memory system according to the PRIOR ART;

 FIG. 2 is a block diagram of a memory system according to the invention;

20 FIG. 3 is a block diagram of an exemplary one of a plurality of cache memory printed circuit boards used in the system of FIG. 2;

 FIG. 4 is a block diagram of an exemplary one of a plurality of front-end directors used in the system of FIG. 2;

25 FIG. 5 is a block diagram of an exemplary one of a plurality of rear-end directors used in the system of FIG. 2;

 FIG. 6 is a block diagram of an exemplary one of a plurality of ASIC control logics used in the cache memories of FIG. 3;

30 FIG. 7 is a block diagram of another embodiment of an interface in accordance with the invention; and

 FIG. 8 is a block diagram of still another embodiment of an interface in accordance with the invention.

- 6 -

Detailed Description

Referring now to FIG. 2, a computer system 100 is shown. The computer system 100 includes a host computer section 112 (e.g., a main frame or open systems computer section) having a plurality of processors, not shown, for processing end-user data. Portions of the processed end-user data are stored in, and retrieved data from, a bank 116 of disk drives through an interface 118. The interface includes a cache memory section 120, here made up of two identical cache memory printed circuit boards 120₀, 120₁, an exemplary one thereof, here memory board 120₀ being shown in, and to be discussed in detail in connection with, FIG. 3. Suffice it to say here, however, that the memory board 120₀ includes an array of DRAMs, here arranged in four memory regions, i.e., memory region A, memory region B, memory region C and memory region D, as shown in FIG. 3 and described in detail in co-pending patent application Serial No. 09/052,268, entitled "Memory System" filed March 31, 1998, inventor John K. Walton, the entire subject matter thereof being incorporated herein by reference.

Referring again to FIG. 2, the interface 118 also includes a plurality of, here eight directors 122₀-122₇. Here, four of the directors, i.e., directors 122₀-122₃, are front-end one of the directors and are coupled to the host computer 112. Here, four of the directors, i.e., directors 122₄-122₇, are rear-end one of the directors and are coupled to the bank of disk drives 116. Each one of the front-end directors 122₀-122₃ is identical in construction, an exemplary one thereof, here front-end director 122₀ being shown in FIG. 4. It is noted that here a cross-bar switch 123 is included to couple each one of a plurality, here four example four, processors, not shown in the host computer 112, to either port P₀ of cache memory 120₀ via serial bus 126_{0,0} or port P₀ of cache

- 7 -

memory 120₁ via serial bus 126_{0,1}, as indicated. Here, each one of the serial busses is a four wire bus having a differential pair of receive wires and a differential pair of transmit wires. Likewise, each one of the rear-end directors 122₄-122₇, is identical in construction, an exemplary one thereof, here rear-end director 122₄, being shown in FIG. 5. It is noted that here a cross-bar switch 123 equivalent to that shown in FIG. 4, is included to couple each one of a plurality, here four example four, rows disk drives, not shown, in bank 116, to either port P₄ of cache memory 120₀ via serial bus 126_{4,0} or port P₄ of cache memory 120₁ via serial bus 126_{4,1}, as indicated.

Referring again to FIG. 2, the interface 118 also includes an interface state data bus section 124, here made up of four interface state data parallel, here 72 wire, busses, i.e., bus A, bus B, bus C, and bus D, for carrying interface state data through the interface 118. The interface state data bus section 124 is coupled to: the front-end directors 122₀-122₃, the rear-end directors 122₄-122₇; and all to the cache memory printed circuit boards 120₀, 120₁. Thus, each one of the four busses is a multi-drop bus. The interface 118 also includes a plurality of, here 16, serial end-user data busses 126_{0,0}-126_{7,1} for carrying end-user data, as indicated. Each one of the plurality of end-user data busses 126_{0,0}-126_{7,1} has a first end coupled to a corresponding one of the plurality of directors 122₀-122₇, and a second end coupled to the memory section 120. More particularly, and considering director 122₀, such director is coupled to memory board 122₀ through end-user data bus 126_{0,0} and to memory board 122₁ through end-user data bus 126_{0,1}. Director 122₁ is coupled to memory board 122₀ through end-user data bus 126_{1,0} and to memory board 122₁ through end-user data bus 126_{1,1}. The other directors are coupled in like manner,

- 8 -

for example, director 122, is coupled to memory board 122, through end-user data bus 126_{7,0} and to memory board 122, through end-user data bus 126_{7,1}, as shown. The plurality of directors 122₀-122, control the end-user data transfer
5 between the host computer 112 and the bank of disk drives 116 through the memory 120 via the end-user data busses 126_{0,0}-126_{7,1} in response to interface state data generated by the directors 122₀-122,. The interface state data is
10 among the directors 122₀-122, through the memory section 120 via the interface state bus section 124.

An exemplary one of the cache memories 120₀, 120₁, here memory 120₀ is shown in detail in FIG. 3. Such memory section 120₀ includes a plurality of, here four
15 random access memory (RAM) regions (i.e. RAM region A, RAM region B, RAM region C and RAM region D, as shown, and a matrix of rows and columns of control logic sections, here Application Specific Integrated circuits (ASICs), i.e. control logic section ASIC A,A ... control
20 logic section ASIC D,D. Each one of the four columns of control logic section ASICs is coupled to a corresponding one of the interface state data busses A, B, C, and D, respectively, as shown. More particularly, a first column of control logic sections (i.e., ASICs A,A; B,A;
25 C,A and D,A) are coupled to the A bus. A second column of control logic sections (i.e., ASICs A,B; B,B; C,B and D,B) are coupled to the B bus. A third column of control logic sections (i.e., ASICs A,C; B,C; C,C and D,C) are coupled to the C bus. A fourth column of control logic
30 sections (i.e., ASICs A,D; B,D; C,D and D,D) are coupled to the D bus.

Each one of the rows of the control logic sections ASIC A,A ... ASIC D,D is coupled to a corresponding one of the four RAM regions, RAM region A ... RAM region D,
35 via a DATA/CHIP SELECT, as indicated. The first row of

- 9 -

ASICs A,A; A,B; A,C; and A,D is coupled to the DATA/CHIP
SELECT BUS of RAM region A. The second row of ASICs B,A;
B,B; B,C; and B,D is coupled to the DATA/CHIP SELECT BUS
of RAM region B. The third row of ASICs C,A; C,B; C,C;
5 and C,D is coupled to the DATA/CHIP SELECT BUS of RAM
region C. The fourth row of ASICs D,A; D,B; D,C; and D,D
is coupled to the DATA/CHIP SELECT BUS of RAM region D.
It should be noted that the control logic sections ASIC
A,A ... ASIC D,D in each of the four rows thereof are
10 interconnected through an arbitration bus, not shown, in
a manner described in detail in co-pending patent
application entitled "Bus Arbitration System", Serial No.
08/996,807, filed December 23, 1997, inventors
Christopher S. MacLellan and John K. Walton, assigned to
15 the same assignee as the present invention, the entire
subject matter thereof being incorporated in this patent
application.

Each one of the rows of the control logic sections
ASIC A,A ... ASIC D,D is coupled to a corresponding one
20 of the four RAM regions, RAM region A ... RAM region D,
via an MEMORY ADDRESS/CONTROL BUS, as indicated. The
first row of ASICs A,A; A,B; A,C; and A,D is coupled to
the MEMORY ADDRESS/CONTROL BUS of RAM region A. The
second row of ASICs B,A; B,B; B,C; and B,D is coupled to
25 the MEMORY ADDRESS/MEMORY BUS of RAM region B. The third
row of ASICs C,A; C,B; C,C; and C,D is coupled to the
MEMORY ADDRESS/ CONTROL BUS of RAM region C. The fourth
row of ASICs D,A; D,B; D,C; and D,D is coupled to the
MEMORY ADDRESS/ CONTROL BUS of RAM region D.

30 The cache memory 120, also includes a coupling
node 130 adapted to couple any one of the ports P_0 - P_7 to
any one of four ports P_A , P_B , P_C , and P_D . Thus, the
coupling node 130 adapted to couple any one of the ports
 P_0 - P_7 to any one of the rows of DATA/CHIP SELECT BUSSES,
35 i.e., any one of the ports P_0 - P_7 to any one of the four

- 10 -

memory regions, A, B, C or D, selectively in response to control signals produced by the ASICs A,A through D,D on a COUPLING NODE CONTROL BUS of each one of the ASICs A,A through D,D. It should be noted that the end-user data
5 is selectively coupled through the coupling node 130 in accordance with routing information fed thereto by the ASICs A,A through D,D in a manner to be described.

The coupling node 130 includes a cross-bar switch section 132 having a plurality of, here two, cross-bar
10 switches 132a, 132b. The coupling node 130 also includes a plurality of, here four, data selectors 134a through 134d. Each one of the cross-bar switches 132a, 132b is a 4x4 cross-bar switch controlled by control signals fed thereto by the ASICs A,A through D,D. Thus, each one of
15 the cross-bar switches 132a, 132b has four input/outputs (here coupled to ports P_0 - P_3 , P_4 - P_7 , respectively, as indicated and four output/inputs coupled to a corresponding one of a pair of input/outputs of the four data selectors 134a through 134d. Each one of such
20 control logic sections ASICs A,A-D,D is identical in construction, an exemplary one thereof, here control logic section ASIC A,A being shown in detail in FIG. 6 to include a control logic 150 having control logic and a buffer memory 152 described in the above-referenced co-
25 pending patent application entitled "TIMING PROTOCOL FOR A DATA STORAGE SYSTEM", inventor John K. Walton, Serial No. 08/996,809, filed December 23, 1997, assigned to the same assignee as the present invention, the entire subject matter thereof being incorporated herein by
30 reference. The ASIC (A,A) controls transfer of data between the buffer memory 152 and the one of the plurality of buses (i.e., A bus, B bus B, C bus and D bus) coupled to the control logic section ASIC A,A, here bus A. The control logic section ASIC A,A is adapted to
35 produce a control/data bus request for the one of the

- 11 -

control/data buses coupled thereto (here RAM region A) and is adapted to effect the transfer in response to a control/data bus grant fed to the control logic section (here ASIC A,A) in accordance with a protocol described
5 in the above-referenced co-pending application entitled "TIMING PROTOCOL FOR A DATA STORAGE SYSTEM", Serial No. 08/996,809, inventor John K. Walton, filed December 23, 1997, the entire subject matter thereof being incorporated herein by reference. The control logic
10 section ASIC A,A also includes a bus arbitration section 153 described in detail in connection with the above referenced patent application entitled "Bus Arbitration System", filed December 23, 1997, inventors Christopher S. MacLellan and John K. Walton. Here, however, the
15 arbitration section also arbitrates for memory contention with one addition bus, the one end-user data serial busses coupled to either port P_A, P_B, P_C, or P_D by the cross-bar switch section 132, in addition to busses A, B., C, and, D. The control logic 150 includes a decoder
20 157 for decoding the eight chip select signals and one read/write signals on the A bus as described in detail in the above referenced co-pending patent application Serial No. 09/052,268. The decoder 157 produces address, control, and clock for the memory region A on the MEMORY
25 ADDRESS/CONTROL BUS, as indicated.

The routing information is fed to the ASICs A,A through D,D via the interface state data busses A, B, C, and D. Thus, for example, if a burst end-user data is to be transferred from director 122, to memory region A of
30 cache memory 120₀, it is first noted that such director 122, is connected to interface state data bus A and therefore the coupling node routing information is for such end-user data transfer is placed by director 122, on interface state data bus A. Referring now also to FIGS.
35 3 and 6, the ASIC A,A control logic, in response to the

- 12 -

coupling node routing information on bus A, presents control information on the COUPLING NODE CONTROL BUS coupled thereto to couple port P_i of the cache memory 120₀ (which is connected to director 122, via end-user bus 126_{0,0}) through cross-bar switch 132a and selector 134a to port P_A.

Referring now to FIG. 7, another embodiment of an interface, here interface 118' is shown. Here, the interface 118' includes 16 directors 122'₀ - 122'₁₅ and four cache memory printed circuit boards 120'₀-120'₃, interconnected through parallel, multi-drop busses A, B, C and D for carrying interface state data and sixty-four sets of serial, point-to-point busses 126'_{0,0}-126'_{0,3} through 126'_{15,0}-126'_{15,3} for carrying end-user data, as shown. It is noted that here there are four cross-bar switches 132 for each printed circuit board 120'₀-120'₃. Further, here each one of the directors 122'₀-122'₁₅ includes four cross-bar switches 132', each one being a 4x4 cross-bar switch.

Referring now to FIG. 8 another embodiment of an interface, here interface 118" is shown. Such interface 118" includes 16 directors 122"₀ - 122"₁₅ and four cache memory printed circuit boards 120"₀-120"₃, interconnected through four parallel, multi-drop busses TH, TL, BH, and BL, such busses being arranged as shown and as described in co-pending patent application entitled "Data Storage System", inventors Daniel Castel, et al., assigned to the same assignee as the present invention, and filed on the same day as this patent application, the entire subject matter thereof being incorporated herein by reference. As described in such co-patent application, the cache memory is arranged to two set; one set having low address and one set having high addresses. Thus, here cache memory printed circuit boards 120"₀ and 120"₁ are included in the high address memory set and cache memory printed circuit

- 13 -

boards 120", and 120", are included in the low address memory set. Here, however, the parallel, multi-drop busses TH (i.e., top high), TL (i.e., top low), BH (bottom high), and BL (bottom low B) are used to carry interface state words and end-user data is carried by serial, point-to-point busses 126"_{0,0} through 126"_{15,3}, as indicated in FIG. 8.

Other embodiments are within the spirit and scope of the appended claims.

10 What is claimed
is:

- 14 -

1. A data storage system wherein end-user data is transferred between a host computer and a bank of disk drives through an interface, such interface, comprising:

a memory;

5 a plurality of directors, at least one front-end one of the directors being in communication with the host computer and at least one rear-end one of the directors being in communication with the bank of disk drives;

10 an interface state data bus section, for carrying interface state data, such interface state data bus section in communication with: both the at least one front-end one and the at least one rear-end one of the directors; and to the memory;

15 a plurality of end-user data busses, for carrying end-user data, each one of the plurality of end-user data busses having a first end coupled to a corresponding one of the plurality of directors and a second end coupled to the memory; and

20 wherein such plurality of directors control the end-user data transfer between the host computer and the bank of disk drives through the memory via the end-user data busses in response to interface state data generated by the directors, such generated interface state data
25 being transferred among the directors through the memory via the interface state bus.

2. The system recited in claim 1 wherein the end-users data busses are serial busses.

3. The system recited in claim 1 wherein the
30 interface state data buss section includes parallel busses.

- 15 -

4. The system recited in claim 3 wherein the parallel busses are coupled to the directors in a multi-drop configuration.

5. The system recited in claim 3 wherein the end-user data busses are serial busses.

6. The system recited in claim 5 wherein the parallel busses are coupled to the directors in a multi-drop configuration.

7. The system recited in claim 1 including a coupling node and wherein each the memory has a plurality of regions and wherein the each one of the end-user data buses is coupled to the plurality of regions selectively through coupling node.

8. The system recited in claim 7 wherein the coupling node includes a cross-bar switch.

9. The system recited in claim 3 wherein the interface state data bus section includes a plurality of parallel busses, each one thereof being coupled to a one of the plurality of directors and to the memory.

10. A method of operating a data storage system wherein end-user data is transferred between a host computer and a bank of disk drives through an interface, such method comprising:

providing a memory;

providing a plurality of directors, at least one front-end one of the directors being in communication with the host computer and at least one rear-end one of the directors being in communication with the bank of disk drives;

- 16 -

providing a plurality of interface state data busses for carrying interface state data, interface state data busses being in communication with: both the at least one front-end one and the at least one rear-end one of the directors; and to the memory;

providing a plurality of end-user data busses, for carrying end-user data, each one of the plurality of end-user data busses having a first end coupled to a corresponding one of the plurality of directors and a second end coupled to the memory; and

wherein such plurality of directors control the end-user data transfer between the host computer and the bank of disk drives through the memory via the end-user data busses in response to interface state data generated by the directors, such generated interface state data being transferred among the directors through the memory via the interface state bus.


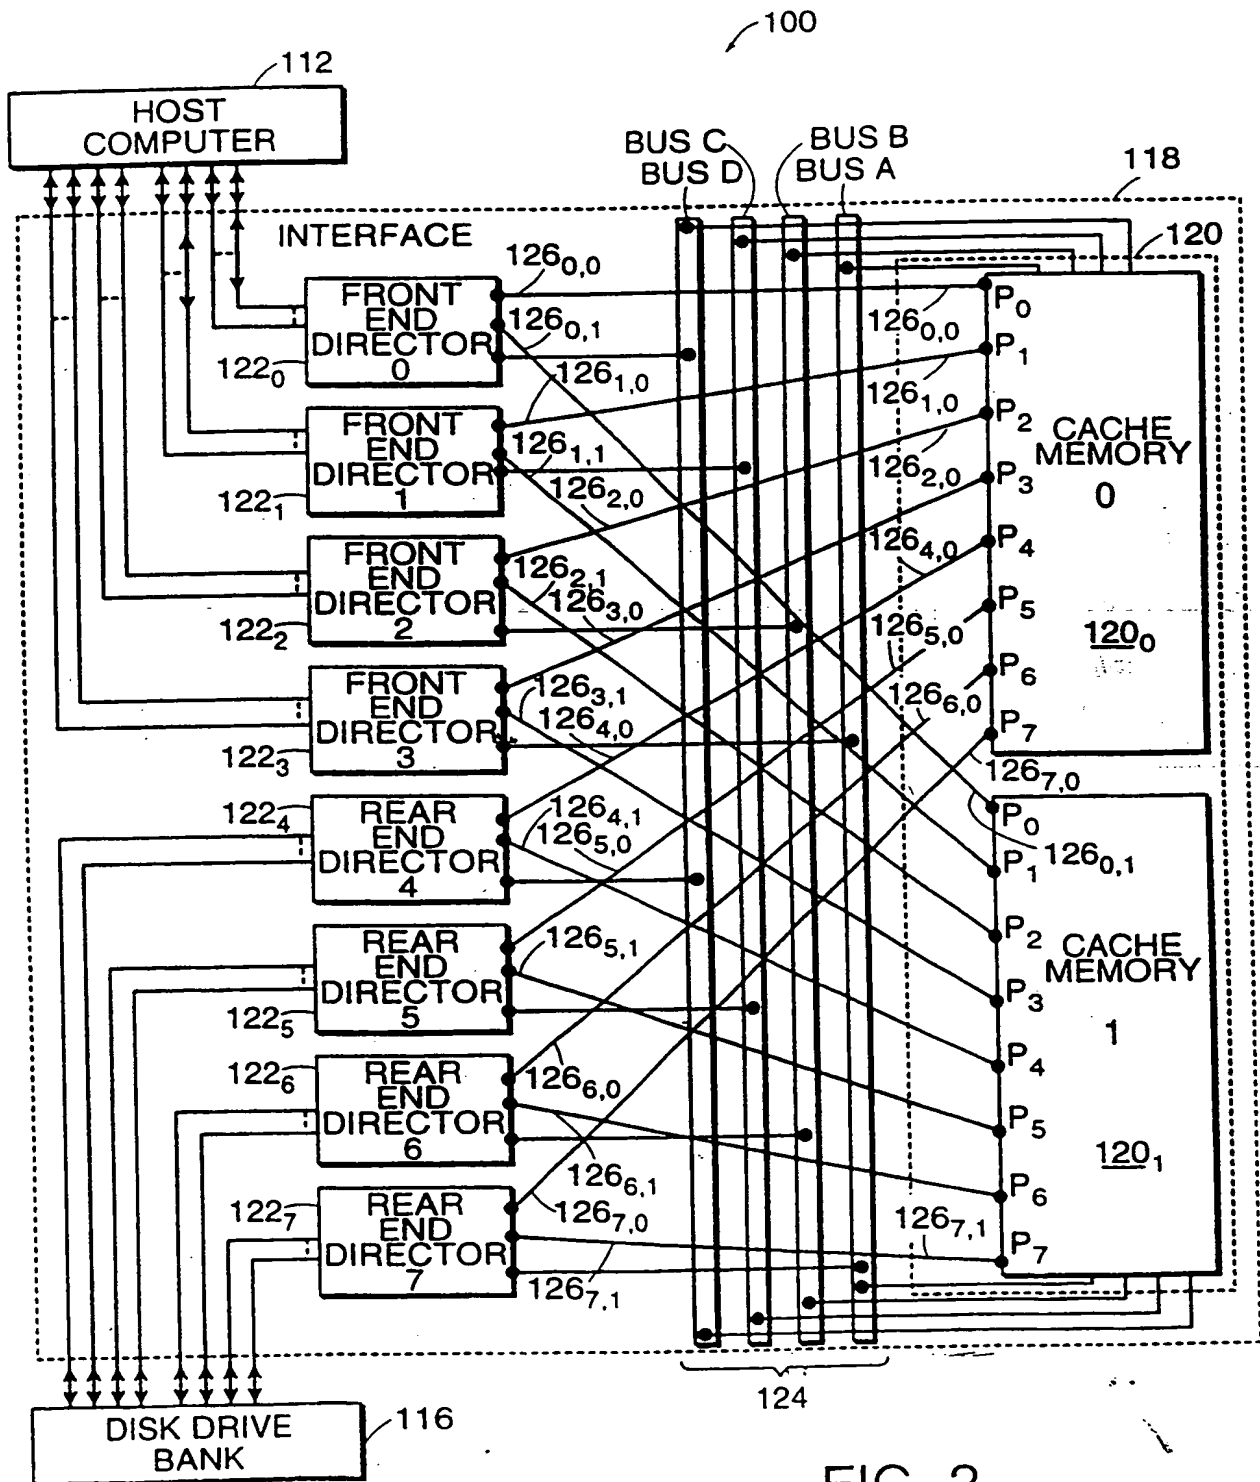


FIG. 1
PRIOR ART



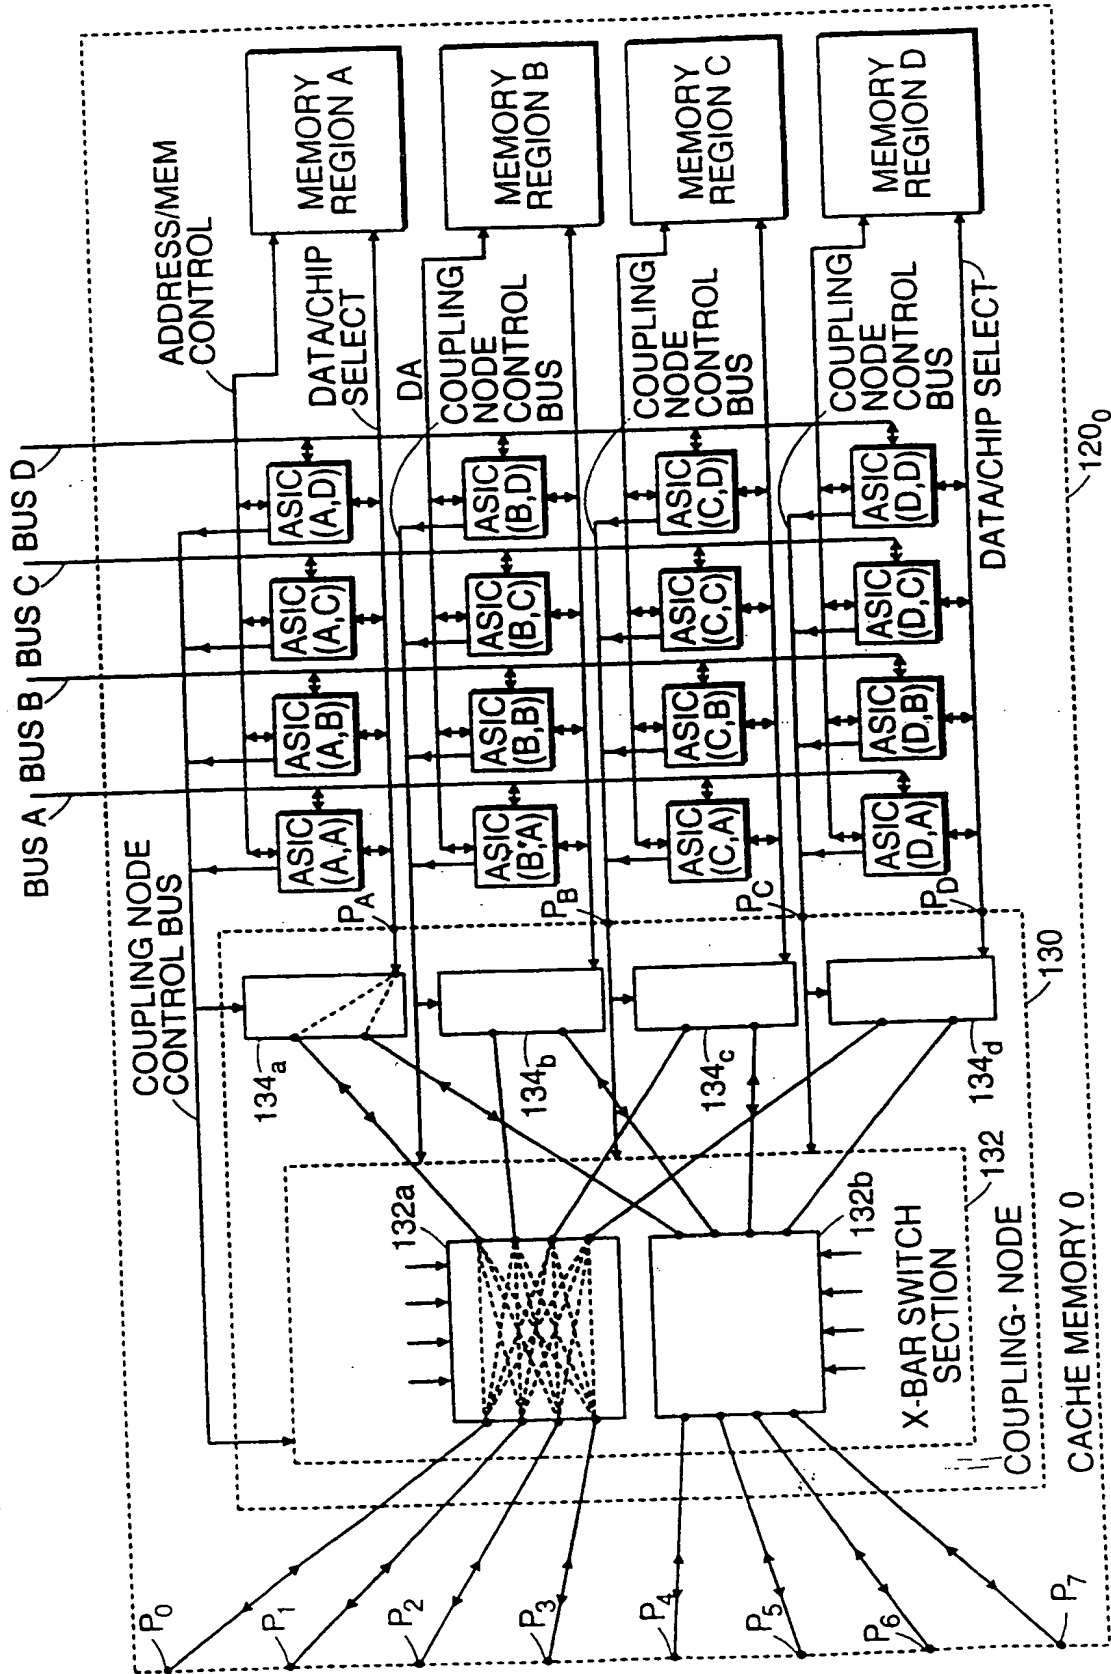


FIG. 3




FIG. 4




FIG. 5




FIG. 6

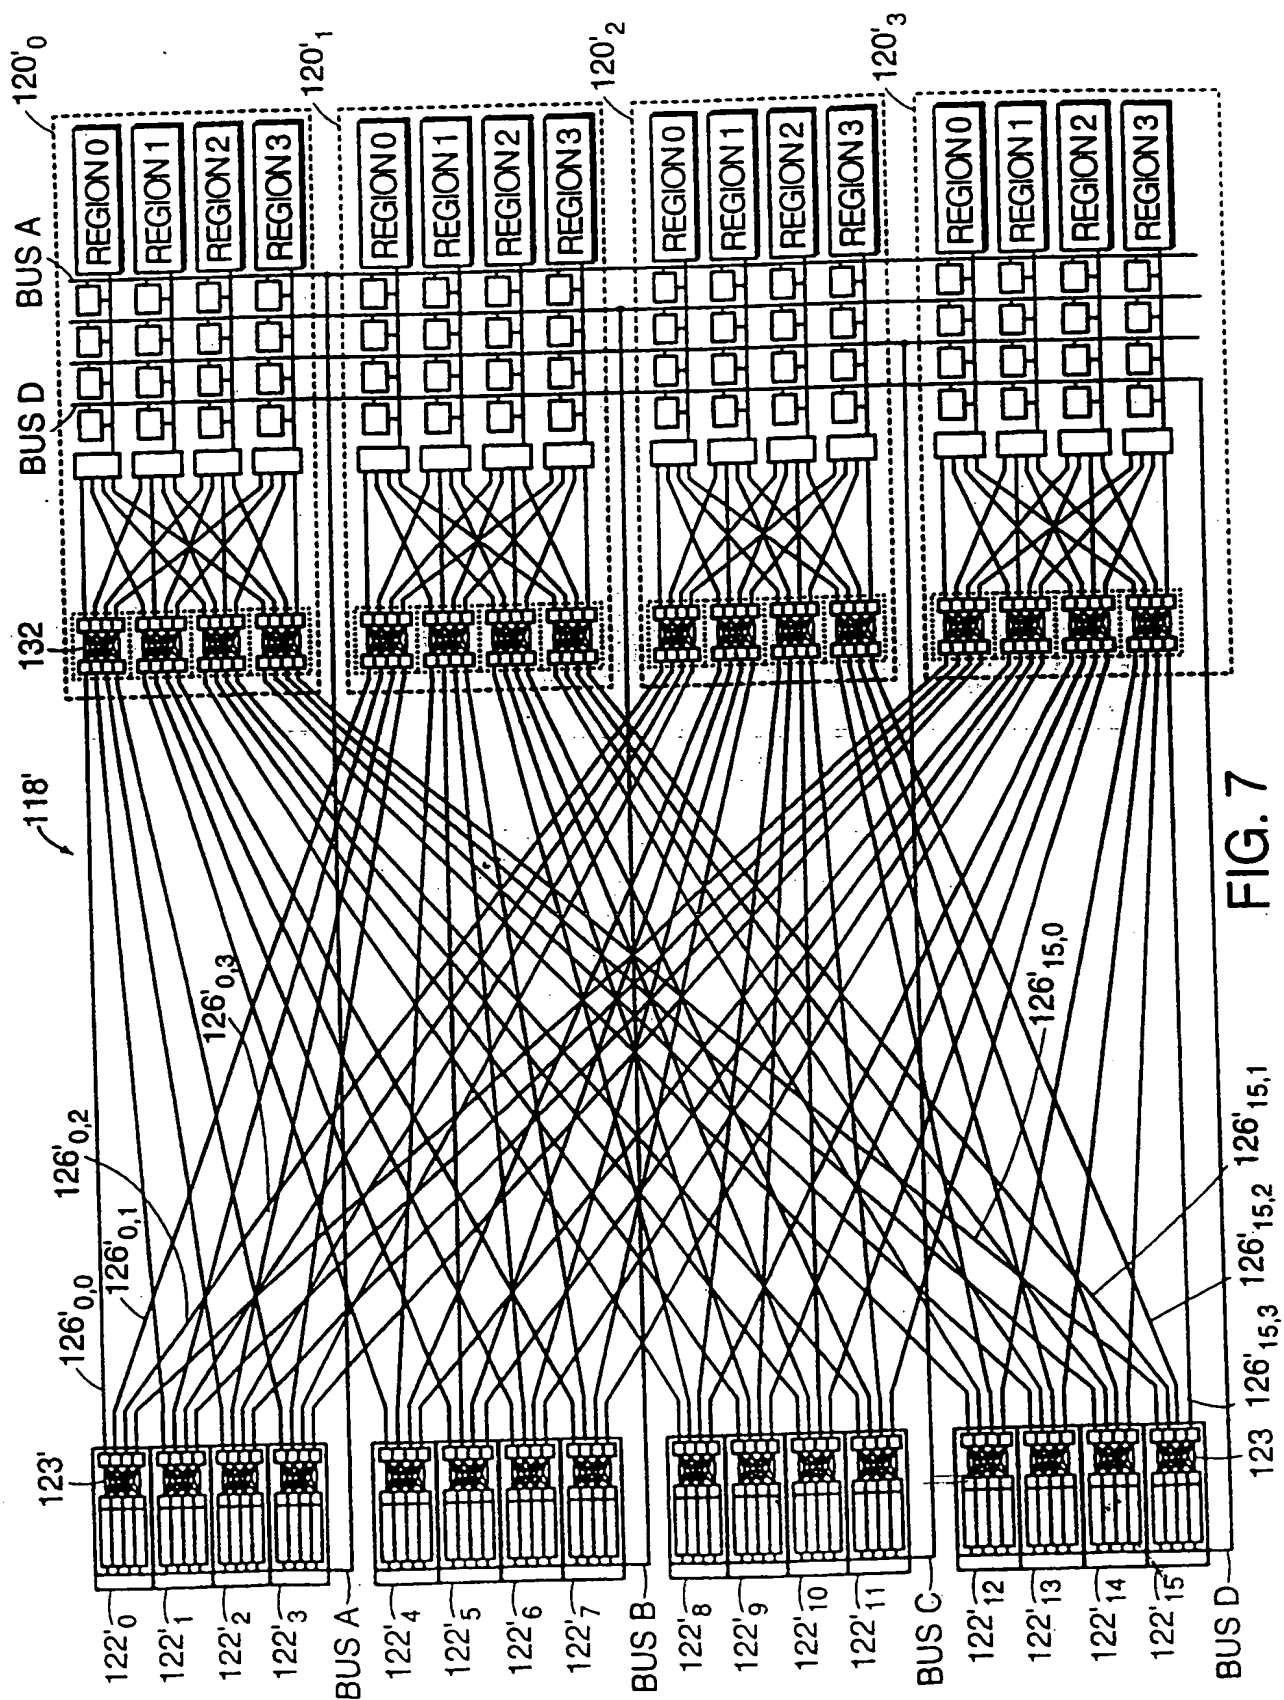


FIG. 7

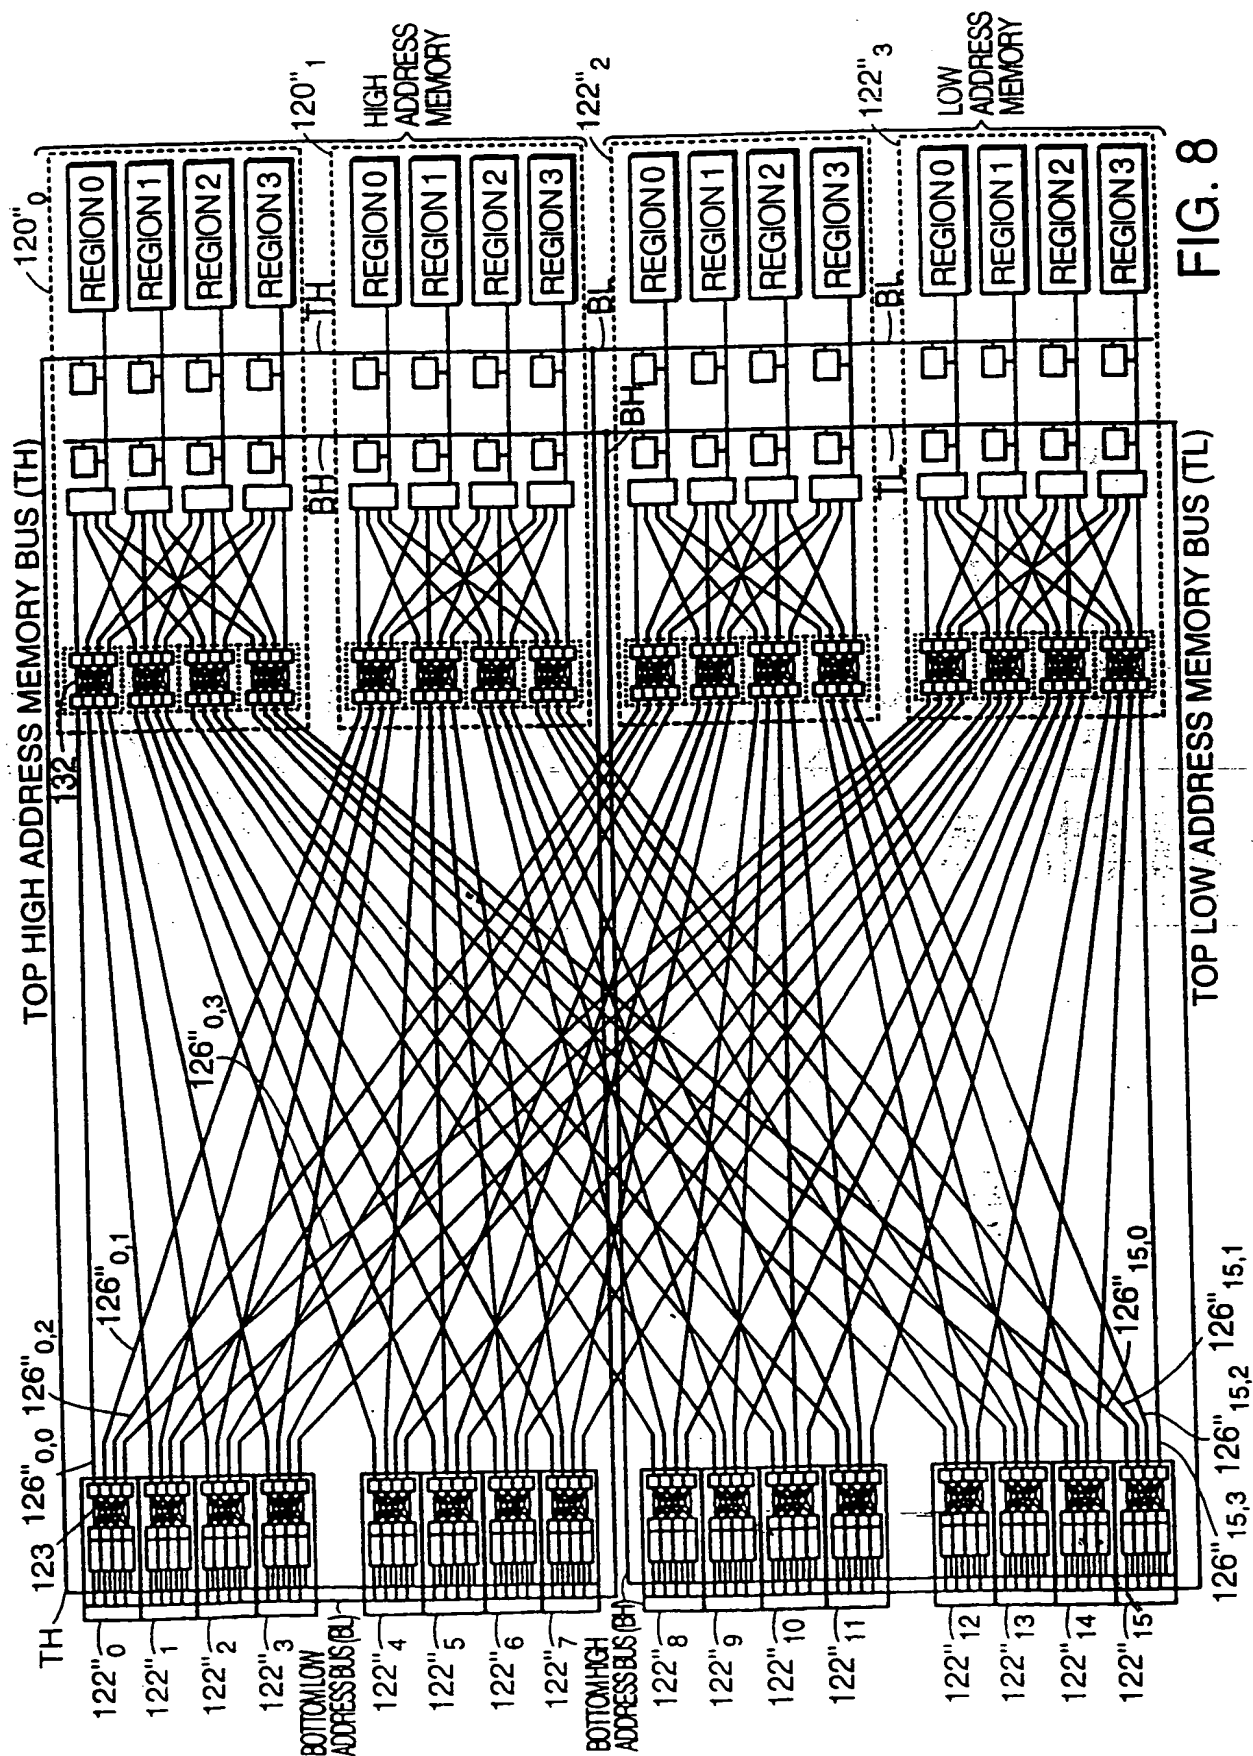


FIG. 8

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 99/31213

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 G06F13/12 G06F12/08 G06F3/06 G06F13/40

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 787 265 A (LESEM) 28 July 1998 (1998-07-28) figure 1	1,10
A	US 5 819 054 A (NINOMIYA ET AL.) 6 October 1998 (1998-10-06) the whole document	1,10
A	EP 0 510 821 A (INTERNATIONAL BUSINESS MACHINES CORPORATION) 28 October 1992 (1992-10-28) abstract	1,10
A	EP 0 795 812 A (HITACHI, LTD) 17 September 1997 (1997-09-17)	

☐

Further documents are listed in the continuation of box C.

☒

Patent family members are listed in annex.

* Special categories of cited documents:

A document defining the general state of the art which is not considered to be of particular relevance

E earlier document but published on or after the international filing date

L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

O document referring to an oral disclosure, use, exhibition or other means

P document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

G document member of the same patent family

Date of the actual completion of the international search

22 May 2000

Date of mailing of the international search report

31/05/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Absalom, R

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 99/31213

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5787265	A	28-07-1998	NONE	
US 5819054	A	06-10-1998	JP 7020994 A US 6012119 A DE 4422786 A	24-01-1995 04-01-2000 12-01-1995
EP 510821	A	28-10-1992	JP 2042421 C JP 4328653 A JP 7076942 B BR 9201347 A CA 2062909 A,C US 5386511 A	09-04-1996 17-11-1992 16-08-1995 01-12-1992 23-10-1992 31-01-1995
EP 795812	A	17-09-1997	JP 8137630 A WO 9615488 A	31-05-1996 23-05-1996